Correlating IDS Alerts with Vulnerability Information

December 2002
(Updated January 2009)

Ron Gula
Chief Technology Officer

Table of Contents

Introduction

This paper will illustrate a variety of the approaches and theories that can be used to correlate intrusion detection system (IDS) logs with vulnerability data. Several models will be presented and their benefits and drawbacks will be discussed. The goal will be to illustrate several methods that vulnerability information can be used to illicit high quality alerts from IDS logs that are primarily false positives.

Intrusion Detection Systems and False Positives

Sources of False Positives

Intrusion detection systems are designed to search network activity (we are considering both host and network IDS detection) for evidence of malicious abuse. When an IDS algorithm "detects" some sort of activity and the activity is not malicious or suspicious, this detection is known as a false positive. It is important to realize that from the IDS's perspective, it is not doing anything incorrect. Its algorithm is not making a mistake. The algorithm is just not perfect. IDS designers make many assumptions about how to detect network attacks.

An example assumption could be to look for extremely long URLs. Typically, a URL may be only 500 bytes long. Telling an IDS to look for URLs longer than 2000 bytes may indicate a denial of service attack. A false positive could result from some complex e-commerce web sites that store a wide variety of information in the URL and exceed 2000 bytes.

Another example would be to look for a reference to the Unix "**`/etc/passwd`**" file in URLs. This file is typically referred to during exploitation of CGI-BIN enabled web servers. An IDS that simply looked for the occurrence of "**`/etc/passwd`**" in any web URL would false positive on someone going to Google and conducting a web search on "**`/etc/passwd`**".

Impact

Because IDS solutions produce false positives, this has several different impacts on how they are deployed and operated on modern networks.

First, the IDS needs to be extensively tuned such that the false positive level is acceptable. Typically, when an IDS is turned on with its default policy, it generates copious amounts of alerts and logs. These logs take up disk space, make for slow database queries and make alerting difficult because there are just too many alerts. This tuning process is iterative, requires extensive knowledge of how the IDS works and extensive knowledge of the monitored network's active protocols and applications. As the network application and usage changes, there exists a chance for large numbers of false positives to be generated again.

Second, because of the false positives, each event must be analyzed by an expert. This analysis usually occurs as an aggregate of all the events once or twice a day and not on a per event basis. For example, an IDS analyst will most likely look at a summary of all unique events detected by their IDS. They would then manually select the alerts that are of interest to them, requesting more information from the IDS user interface. The analyst is using their experience to find the IDS events that are high-quality. In some cases, if a

pattern of false positives or source of false positives can be identified, they can be filtered out with a change to the IDS's policy.

Third, because of the false positives, the real-time source of IDS events is not used to drive automated response systems such as reconfiguration of a firewall, after hours alerting of the security staff or automatically filling of a trouble ticketing system. If this sort of automated notification is used, there is typically extremely heavy filtering. This limits any sort of network outage that can be caused from an IDS reconfiguring a router, switch or firewall.

Vulnerability Correlation with IDS Alerts

Immediate Benefits

Many of the false positives associated with an IDS can be mitigated by considering the vulnerabilities of the protected network. At a high level, if an IDS knows that a system is vulnerable to a particular vulnerability, then it should only concern itself with attacks against that particular vulnerability. We will discuss the merits of continuing to detect attacks for which we know a system to not be vulnerable later in this paper.

If such a system existed, then we can expect high quality alerts to be generated. The system knows what the vulnerabilities are and it knows that a particular vulnerability is being exploited. This level of information results in a higher level of confidence that a system is under immediate threat. Because of this correlation, better decisions can be made in an automated fashion. These include firewall rule changes to drop the attacker and mailing the security staff to notify them in real-time of the attack.

Limitations

Such a system is not perfect though. Several cases exist where the data generated by the system is subject to further analysis. At the root of these limitations are the concept of false positives and negatives within both the IDS and the vulnerability assessment (VA) portions of the system. We already discussed false positives for an IDS. A false negative is when the IDS is presented with a valid attack and it is not detected. Similarly with VA technology, it can false positive on the detection of a vulnerability and it can also false negative and miss a vulnerability. The limitations of a system that correlates IDS and VA data to reduce false positives can roughly be categorized in nine categories that correlate false positive, false negative and true IDS alerts with false positive, false negative and true vulnerabilities. These categories are identified and discussed individually below and guidelines used to minimize or maximize their impact are also presented.

1. False Positive IDS Alerts with False Positive Vulnerabilities

With this scenario, an IDS has detected an attack against a system that is not a real attackand the VA has detected the corresponding vulnerability as well, even though the system is not vulnerable at all. This is a worse case scenario because our system will generate a supposedly "high quality" alert, even though the alert is false. To minimize this from occurring, the best approach is having a tuned IDS and a tuned VA solution. When this scenario occurs, it should be immediately corrected for in the IDS and VA solutions.

Most false positive alerts are generated by normal network traffic and should be generated often. Because of this, scenarios of this type should be tuned out over time. Having said that, the most likely situation for this scenario is when a VA solution and IDS solution are

automatically updated with new checks simultaneously and both checks produce false positives.

This scenario can be identified by looking at the targeted OS, application, IDS event and vulnerability and seeing each of these pieces of information not line up.

2. False Positive IDS Alerts with False Negative Vulnerabilities

This is a non-issue. A false positive IDS alert going to a system that happens to be vulnerable to the alert is tempting to consider as more valuable, but by-definition, this scenario states that the IDS event is a false positive. Since no vulnerability is registered, the system will ignore this scenario.

3. False Positive IDS Alerts with True Vulnerabilities

This is the desired scenario. When the true state of a vulnerability is known, any IDS alert can be thrown away if it is not applicable. In this case, most of the IDS alerts will be thrown away as they are not targeting a vulnerability. There is a chance though, that there will be a false positive IDS event that has targeted a valid vulnerability. This is the similar to scenario 1 in that our system has generated a high quality alert, which is in fact a false positive.

4., 5., 6. False Negative IDS Alerts with False Positive, False Negative and True Vulnerabilities

This series of scenarios is very difficult to apply VA/IDS correlation to. If an IDS can be bypassed, then there is no IDS event to correlate in the first place. Two of our models for VA/IDS correlation will attempt to compensate for this, but the reality is if an IDS can be blinded, then the entire system can be bypassed.

7. True IDS Alerts with False Positive Vulnerability

When an IDS detects a “real” attack and the VA system has incorrectly said that the targeted system is also vulnerable to the attack, then our system will generate a high quality alert, which may actually not be that high quality. This is just as serious as scenario 1 because a high quality event has been generated when in fact this attack has little chance of succeeding.

This scenario can be identified when the actual IDS event is investigated and the targeted systems are not vulnerable to the attack. They will most likely result in seeing that the vulnerabilities reported are off the mark. For example, saying that there is a vulnerability in the Sendmail server at a particular IP address that is not running Sendmail, but another mail transfer agent.

8. True IDS Alerts with False Negative Vulnerability

In this scenario, the IDS has done its job correctly, but the VA system has said we are not vulnerable to this attack, so we can ignore it. This is another worse case scenario in that we threw away some good IDS data because we had bad vulnerability data. One of our models discussed later can mitigate this scenario from occurring. Short of building better VA technology, there is not much more that can be done with this scenario.

Recognizing this scenario may occur when the IDS analysts detect a true break in and the targeted vulnerability was found to not be detected by the VA solution. This is most likely

going to occur in a zero-day situation in which a server is compromised with a vulnerability that has yet to be published.

9. True IDS Alerts with True Vulnerabilities

This is also a desired solution in which everything works correctly. In this case, the IDS observes the attack and the VA system correctly identified the targeted system as being vulnerable to the specific form of attack.

Correlation

The most likely way systems will be built to correlate IDS and VA data will be either through proprietary solutions or open standards.

For proprietary solutions, we are talking about single vendor solutions that offer separate VA and IDS products. In this case, the vendor has assigned unique identification values to each of their VA checks and IDS signatures. Correlating the two is done with a direct lookup. For example, an IDS check that looks for the "Anonymous FTP" exploit will likely have a reference to the vendor's VA check for "Anonymous FTP" as well.

When a single source for VA and IDS technology is not available, loose correlation of open standards is likely. There exist a number of "open" standards or references that can be used to correlate disparate VA and IDS solutions. These include Mitre's Common Vulnerability Enumeration program (http://cve.mitre.org/), the Computer Emergency Response Team's CERT advisories (http://www.cert.org/), Bugtraq (http://www.securityfocus.com/bid) and in some cases, Nessus vulnerability IDs can be used.

To do this, a solution would have to look at all of the vulnerabilities they check for and extract relevant CVE, CERT or other information associated with each vulnerability check. They would then have to associate which of their monitored systems had particular vulnerabilities. Then, they would have to repeat the process for each of the IDS devices they used to correlate logs from. This can be difficult, as many IDS vendors do not directly publish their signature knowledge base in a form that can be utilized for this purpose. And finally, whenever an IDS event occurred, a quick check that was able to tie the alert to a specific vulnerability and then to a specific vulnerable system is required.

It should also be noted that an IDS can detect a lot more activity than is directly traceable to a vulnerability. For example, a network probe like a ping sweep or an OS identification event can indicate a potential intrusion, but does not directly correlate to a known vulnerability.

VA / IDS Correlation Models

Several approaches can be used to leverage knowledge of vulnerabilities on a network system and limit the number of false positives generated by an IDS. Each of these models is explained and then has its strengths and weaknesses analyzed.

Vulnerability Based IDS Policies

In this model, knowledge of the vulnerabilities on a network (and perhaps even the topology) is used to create policies for each of the IDS devices. The end result is that each IDS device will only be looking for network activity that is known to be vulnerable on a

particular network. How much information you give an IDS can be an asset and a burden. For example, if we knew that none of our protected systems were vulnerable to Telnet buffer overflow attacks, then it would be very useful to disable all of the signatures on all of our NIDS devices. On the other hand, attempting to store a Class B (60,000+ IP addresses) of vulnerability information in such a way that can be used by a real-time 100 MB NIDS device, is something that has not been achieved by the IDS industry.

The strength of this solution is that your IDS sensors may run quicker, if they can handle the information from the VA. Running with less signatures is a great way to increase the speed of any host or network IDS. Another benefit is that the alerts generated are already high quality. It follows that there would be less false positives as well.

The weakness of this solution is that it may throw away interesting attack information. To date, an IDS system has not been constructed that could mark each of its IDS alerts if it was vulnerable or not to the IDS attack detected. This would be desirable. Because of this, IDS solutions in this area tend to ignore real attacks that do not target a specific vulnerability. This is sort of like a robber coming to your house and trying your front door, your side door, your patio door and your cellar window. He has not got in yet, but this may be activity you want to see. Unfortunately, on the Internet, this analogy does not hold. On the Internet, this analogy would include someone trying all of your houses entrances several times a day. Having said that, the IDS community is split between people who would rather not be bothered with the enormous amount of information that IDS devices generate and people who wish to look at all of the logs in context with each other. Any system that threw away information will be desirable to people in the first group and not at all wanted by people in the second group.

This model helps with scenarios 1, 2 and 3 because they tend to look for less attacks in general. Compared to a NIDS device that may look for thousands of exploits, a NIDS that was tuned to only look for a handful of exploits will have much less false positives.

As NIDS technologies evolve, it may be possible to configure a higher end solution that supported virtual NIDS devices with this form of vulnerability based policy. Several hardware based NIDS vendors support virtual NIDS engines. That is, they can run more than one instance of their NIDS engine. This usually occurs without a performance penalty. Because of this, a NIDS with this feature could be configured to operate with a specific policy that reflected a robust set of vulnerability information. This can lead to high quality events being generated by the virtual NIDS.

Persistent VA/IDS Correlation

This system maintains a database of network vulnerabilities and correlates it with IDS alerts. When a correlation is found, it can send a high quality alert. Such a system would have to have knowledge of the vulnerability checks performed and have a method to link a large number of the IDS alert types directly with the vulnerability checks. The system should also have a method to query the database of vulnerabilities quick enough to keep up with the feed of IDS events.

This system has the benefit of keeping all of the IDS events that occur and marking the ones that directly correlate to a vulnerability. This allows user interfaces to include a button that could magically remove all of the IDS events, except the ones that targeted a vulnerability. If the reader is familiar with the enterprise security manager (ESM) or security information manager (SIM) product space, another benefit of this technology is to receive IDS events from different solutions. It is very likely that different IDS technologies will

detect attacks in different ways. More information increases the chances of scenario 9 occurring, which finds a valid IDS event targeting a known valid vulnerability.

This system has some drawbacks though. First, not all IDS solutions will directly correlate with each vulnerability check. Knowing what the limitations are and keeping the nine scenarios in mind will help you evaluate the effectiveness of your VA/IDS correlation solution. Second, if the database is out of date, there exists a much higher chance for scenarios 7 and 8 to occur. A network VA scan could be out of date seconds after the scan is completed. However, for most server farms, their configuration does not change that quickly and most scans can be set to occur on a daily basis.

Near-time VA/IDS Correlation

This model is very similar to the above model, but does not maintain a permanent database of vulnerabilities. Instead, as IDS events occur, the network is actively queried for vulnerability information. Typically, the operating system is detected, then the application and then the vulnerability itself. Along the way, if any vulnerability information is derived that does not correlate with the IDS event, the event is discarded.

The advantage of this system over the above system is that a large database of vulnerabilities does not have to be maintained. However, the system is at a disadvantage in high speed IDS environments, where IDS alerts occur at such a rate, that there is no time to stop and check the network to verify. Also, this type of system may present an attacker with an opportunity to conduct a denial of service attack by spoofing some simple attacks in large numbers and having the VA/IDS correlation system (verification system) launch a much more in-depth query against the targeted systems.

Real-time VA/NIDS Correlation

Finally, our last system considers the “ultimate” NIDS in which vulnerability information is derived in real-time. This is different than our first model in which vulnerability information is directly given to the IDS. In this example, if the NIDS can derive a vulnerability (with a signature) they may be able to conduct some VA/IDS correlation in real-time. The complexities of this may not bear much fruit though, as many vulnerabilities cannot be derived passively and must require active interaction by a VA solution. For example, a Windows IIS web server may be vulnerable to many different attacks. As it is patched for these attacks, its banner information does not change. This banner information is what a NIDS would use to determine vulnerability information.

If such a system could be developed and be reliable, accurate and fast enough for use at network speeds, then the alerts that would be generated would be of high value. At a minimum, this solution may be able to automatically throw away checks that are not relevant. In the above example, the NIDS may not know which vulnerabilities the IIS server had, but it would know that an IIS server existed and as such, should not check for IDS exploits other than IIS exploits.

Conclusion

There are many ways to correlate IDS information with vulnerability data. Each of these ways has a variety of benefits and drawbacks. Each way is also susceptible to one of the nine scenarios that exist when correlating this sort of information. Understanding the benefits and drawbacks of any VA/IDS deployment is necessary to increasing the

effectiveness of your security monitoring. In many cases, high quality IDS alerts can be obtained through the use of VA/IDS correlation. This allows for greater automation to take action in real-time against intruders.

About the Author

Ron Gula is a founder and Chief Technology Officer of Tenable Network Security, Inc. Tenable is a company that produces the Lightning Proxy for high-speed Nessus vulnerability scans and the Security Center for correlating IDS data with vulnerability data and making it available to multiple people in multiple organizations. Previously, Mr. Gula was the original author of the Dragon IDS and CTO of Network Security Wizards that was acquired by Enterasys Networks. At Enterasys, Mr. Gula was Vice President of IDS Products and worked with many top financial, government, security service providers and commercial companies to help deploy and monitor large IDS installations. Mr. Gula was also the Director of Risk Mitigation for US Internetworking and was responsible for intrusion detection and vulnerability detection for one of the first application service providers. Mr. Gula worked for BBN and GTE Internetworking where he conducted security assessments as a consultant, helped to develop one of the first commercial network honeypots and helped develop security policies for large carrier-class networks. Mr. Gula began his career in information security while working at the National Security Agency conducting penetration tests of government networks and performing advanced vulnerability research.

About Tenable Network Security

Tenable, headquartered in Columbia, Md., USA, is the world leader in Unified Security Monitoring. Tenable provides agent-less solutions for continuous monitoring of vulnerabilities, configurations, data leakage, log analysis and compromise detection. For more information, please visit us at http://www.tenablesecurity.com/.

TENABLE Network Security, Inc.
7063 Columbia Gateway Drive
Suite 100
Columbia, MD 21046
TEL: 410-872-0555
http://www.tenablesecurity.com/